Dynamic Decision Modeling in Medicine: A Critique of Existing Formalisms

Tze-Yun Leong

Clinical Decision Making Group
MIT Laboratory for Computer Science
Cambridge, MA 02139
(E-mail: leong@lcs.mit.edu)

Dynamic decision models are frameworks for modeling and solving decision problems that take into explicit account the effects of time. These formalisms are based on structural and semantical extensions of conventional decision models, e.g., decision trees and influence diagrams, with the mathematical definitions of finite-state semi-Markov processes. This paper identifies the common theoretical basis of existing dynamic decision modeling formalisms, and compares and contrasts their applicability and efficiency. It also argues that a subclass of such dynamic decision problems can be formulated and solved more effectively with non-graphical techniques. Some insights gained from this exercise on automating the dynamic decision making process are summarized.

INTRODUCTION

Decision analysis is a normative framework for decision making under uncertainty. By systematically formulating, evaluating, and analyzing a graphical *decision model*, the decision analytic approach helps in both gaining better insights into, as well as deriving a set of optimal decisions for the problem at hand. In recent years, some new decision modeling formalisms have been devised to deal with dynamic or sequential decision problems, *i.e.*, decision problems that take into explicit account the effects of time. For instance, a common clinical decision is to choose a course of treatments with efficacies that may vary over time. These new formalisms, *e.g.*, dynamic influence diagrams [12], Markov cycle trees [1] [8], and stochastic trees [6], are based on structural and semantical extensions of conventional decision models such as influence diagrams and decision trees, with the mathematical definitions of semi-Markov processes.

The different modeling formalisms have different pros and cons. There are a lot of controversies and uncertainties about the "best" framework available. Judging from the literature, researchers adopting one framework are usually not familiar with the capabilities and limitations of the others. At times there are attempts at pushing a particular framework to accommodate features that are hard to be incorporated; cumbersome extensions and restrictive assumptions might result.

This paper attempts to unify the existing dynamic decision modeling formalisms by identifying their common theoretical basis, and to distinguish them by classifying the decision problems they can effectively handle. A common basis for comparing the formalisms results from this exercise. This could help researchers decide if a particular framework is suitable for the problem at hand, if reasonable extensions can be made of the framework, and if necessary assumptions can be made of the problem. The quality of the resulting models from different frameworks can then also be compared on a consistent set of metrics.

AN EXAMPLE DECISION PROBLEM

To illustrate the different formalisms, we examine a typical dynamic decision problem in the management of chronic ischemic heart disease (CIHD). The problem, adapted and simplified from [14], is to determine the relative efficacies of different treatments for chronic stable angina (chest pain), the major manifestation of CIHD. The alternatives considered are medical treatments, percutaneous transluminal angioplasty (PTCA), and coronary artery bypass graft (CABG). The major characteristics of this problem are as follows:

Progressive symptoms: All symptoms are assumed to be progressive. If the angina worsens after a treatment, the treatment is likely to have failed, and subsequent actions must be considered.

Recurrent prognosis: Restenosis, i.e., renewed occlusion of the coronary arteries, is possible even after successful treatment.

Sequential decisions: Due to the possibilities of ineffective treatments and restenosis, a sequence of decisions must be made.

Time-dependent efficacies and complications: The efficacies of the treatments in lowering mortality decline as time progresses, partly due to the deteriorating status of the patient. Similarly, the complications of the treatments worsen as time progresses. A major complication for PTCA is perioperative myocardial infarction (MI), or heart attack, which would render an emergency CABG necessary.

Value function: The effectiveness of the different treatments is evaluated with respect to a value function, *e.g.*, quality-adjusted life expectancy (QALE).

DYNAMIC DECISION MODELING

A dynamic decision model is based on a graphical modeling language for explicitly displaying the relevant variables in a dynamic decision problem. In general, such a model consists of the following six components, the first five of which constitute a conventional decision model:

• A set of *decision nodes* listing the alternative actions that the decision maker can take, for instance, the choices of medical therapy, PTCA, and CABG;
• A set of *chance nodes* outlining the possible outcomes or happenings that the decision maker has no control over, for example, the physical status of the patient, the prognostic outcomes of PTCA, *etc.*;
• A single or a set of *value functions* capturing the desirability, in terms of factors like cost, life-expectancy, *etc.*, of each outcome or action;
• A set of *probabilistic dependencies* depicting how the outcomes of each chance node depend on other outcomes or actions;
• A set of *informational dependencies* indicating the information available when the decision maker makes a decision; and
• An underlying *semi-Markov process* governing the evolution in time for the above five components.

We shall now briefly examine three simplified dynamic decision models for the example problem presented earlier. Figure 1 depicts the state transition diagram of the embedded Markov chain for this problem. The states in the chain represent the possible physical conditions or health outcomes of a patient, given any particular treatment $k \in \Omega = \{MedRx, PTCA, CABG\}$. For ease of analysis, assume that each state variable $x \in X =$ {"Well", "Restenosis", "MI", "MI+Restenosis", "Dead"} in the Markov chain is a function of a set of binary health outcome variables $O = \{Status, MI, Restonosis\}$, *e.g.*, "Well":= (*Status = alive, MI = absent, Restenosis = absent*), "MI":= (*Status = alive, MI = present, Restenosis = absent*), "Dead":= (*Status = dead, MI = present or absent, Restenosis = present or absent*), *etc.*

The links in Figure 1 represent possible transitions from one state to another; given any treatment, a *transition probability* is usually associated with each link. The time taken for a transition to take place can be measured in discrete or continuous scale. The formal definition of a Markov chain will be presented later in this paper.

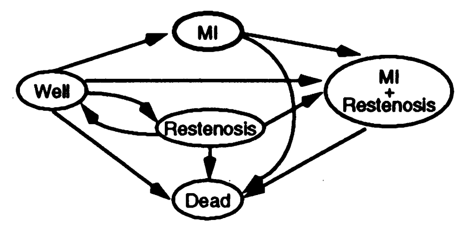

Figure 1: An embedded Makov chain for the example.

In the illustrations of dynamic decision models to follow, a rectangle represents a decision node and an oval a chance node, unless otherwise specified.

Dynamic Influence Diagrams

Figure 2 shows a dynamic influence diagram for the example problem. The shaded ovals represent the value nodes. The number at the end of each node indicates the time period in which the decision/event/value is considered. Each time period corresponds to one transition cycle time in the underlying discrete-time Markov chain; the cycle time may be in any unit and is usually constant within the model. The arcs leading into the chance and value nodes indicate probabilistic dependencies, and the arcs leading into the decision nodes indicate informational dependencies. Embedded in each chance node or value node is a list of the possible values or outcomes of the node, *e.g., "alive"* or *"dead"* for *Status*, and a table of probabilities conditional on its probabilistic predecessors; embedded in each decision node is a list of the alternate treatments and a list of its informational predecessors.

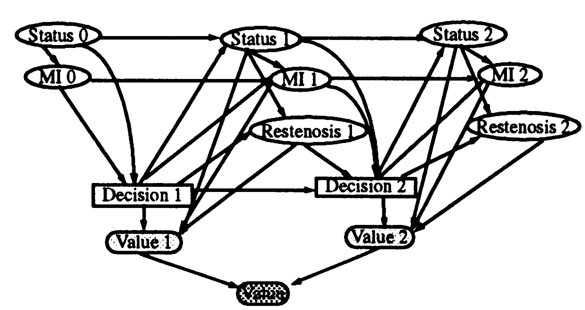

Figure 2: A dynamic influence diagram.

Markov Cycle Trees

Figure 3 depicts a Markov cycle tree [8] for the example problem*. The convention for the notation is such that there is one Markov cycle tree modeling the prognosis of each alternate treatment or strategy in the conventional decision tree, represented by the single decision node in the figure. All the possible combinations of events that could happen between any two transitions of the underlying discrete-time Markov chain, are represented in between the root and the leaves of a cycle tree. The transition cycle time is again constant in a model. The arcs joining the nodes indicate probabilistic dependencies among the nodes. Probabilistic independencies, however, are not explicitly shown as in a dynamic influence diagram.

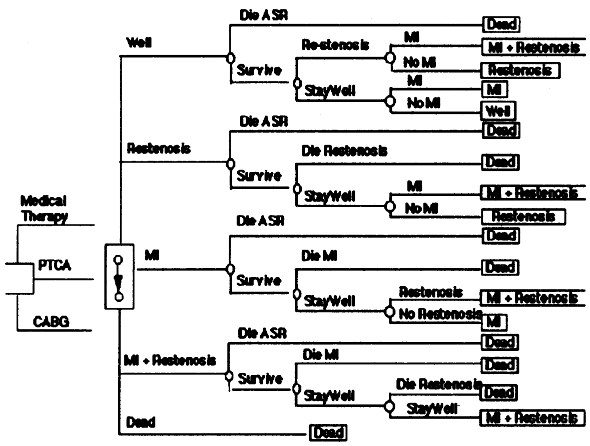

Figure 3: A Markov cycle tree.The symbol denotes the root of the cycle tree, and each $\boxed{<state>}$ a leaf, *i.e.*, the state to start in the next cycle. The label DieASR means die of all other, including age- and sex-related, *i.e.*, natural, causes.

Stochastic Trees

Figure 4 depicts a stochastic tree [6] for the example problem. A stochastic tree, like the Markov cycle tree, is used together with a conventional decision tree to model the prognosis of a decision. Unlike the other two models, however, state transitions in a stochastic tree can occur any time according to some exponential distributions with specific rates, *i.e.*, it models a continuous-time Markov chain. With reference to the embedded Markov chain in Figure 1, only transitions to immediate states are depicted in the stochastic tree. Transition rates, instead of transition probabilities, can be used directly for specifying the probabilistic characteristics of the tree.

*The diagram does not depict sequential decisions. The explanation will be presented later.

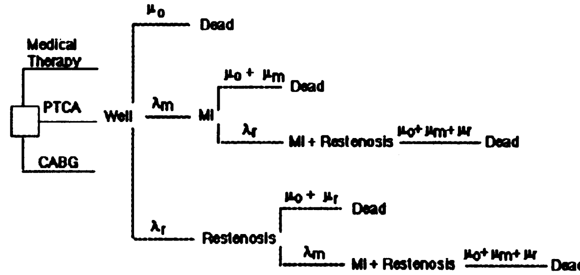

Figure 4: A stochastic tree. Parameters on the arcs indicate transition rates among the relevant states.

SEMI-MARKOV DECISION PROCESSES

The structure and semantics of the models described so far can be more clearly understood in terms of the theoretical basis of dynamic decision modeling: *semi-Markov decision processes*. Formally, a semi-Markov decision process is characterized by the following components [4] [7]:

- A set of n alternate decisions $\Omega = \{1, 2, \ldots, n\}$; and
- A *semi-Markov reward process* denoted by a set of random variables $\{X(t); t \in T\}$, where $X(t) \in \{0, 1, 2, \ldots\}$ is the state of the process at time t, and T is the time-index set, with:

1. an embedded Markov chain denoted by a set of random variables $\{X_m; m \geq 0\}$; such that $X_m = X(T_m)$, where $T_1 < T_2 < T_3 < \ldots$ are the random variables denoting the successive epochs (*i.e.*, instants of time) at which the process makes transitions;
2. n sets of transition probabilities $\{P_{ij}^{(k)}; i \geq 0, j \geq 0, 1 \leq k \leq n\}$ among the states of the embedded chain such that for any given decision $k \in \Omega$, and states *i, j*:

$$P_{ij}^{(k)} = P\{X_{m+1} = j | X_m = i\}$$
$$= P\{X(T_{m+1}) = j | X(T_m) = i\}$$

which also satisfies the *Markovian property*:

$$P_{ij}^{(k)} = P\{X_{m+1} = j | X_m = i\}$$
$$= P\{X_{m+1} = j | X_m = i, X_{m-1} = h, \ldots\}$$

3. n sets of holding times $\{\tau_{ij}^{(k)}; i \geq 0, j \geq 0, 1 \leq k \leq n\}$ among the states of the embedded chain, which are random numbers with corresponding distributions $\{F_{ij}^{(k)}(t); i \geq 0, j \geq 0, 1 \leq k \leq n\}$, such that for any given decision $k \in \Omega$, and states *i,j:*

$$F_{ij}^{(k)}(t) = P\{T_{m+1} - T_m \leq t | X_m = i, X_{m+1} = j\}$$

and

4. n sets of rewards/values $\{r_i^{(k)}(t); i \geq 0, 1 \leq k \leq n\}$ associated with the states of the embedded chain such that for any given decision $k \in \Omega$, $r_i^{(k)}(t)$ is the expected cumulative value achievable in state i of the chain at time t.

A dynamic decision problem can be expressed as the dynamic programming equation, or Bellman optimality equation, of a semi-Markov decision process. For instance, the equation corresponding to the set of discrete time problems that can be captured by the dynamic decision models described earlier is:

$$V_i^*(t) = max_k \{r_i^{(k)}(t) + \sum_j P_{ij}^{(k)} V_i^*(t+1)\}; \quad t \geq 0; i, j \geq 0; 1 \leq k \leq n \quad \text{(EQ 1)}$$

The solution to such an equation is an optimal *policy*, *i.e.*, a sequence of decisions over time t (which could be one single decision repeated indefinitely) that maximizes $V_s^*(0)$, the optimal expected value or reward for a "start" state s, *e.g.*, the "well" state for the example problem, at time 0.

In our example, the set of decisions is $\Omega = \{MedRx, PTCA, CABG\}$. The semi-Markov reward process, with time index set $T \subseteq \{0, 1, 2, \ldots\}$ for the dynamic influence diagram and the Markov cycle tree, and $T \subseteq [0, +\infty]$ for the stochastic tree, is defined by: 1) the embedded Markov chain with state-space X = {"Well", "Restenosis", "MI", "MI+Restenosis", "Dead"} as illustrated in Figure 1; 2) three sets of transition probabilities among the states in X, corresponding to the decisions in Ω; 3) constant holding times with distributions $F_{ij}^{(k)}(t) = 1(t-1)$, where $1(t-1)$ is a step function at time $t = 1$ (in any unit) for the dynamic influence diagram and the Markov cycle tree, and exponential holding times with distributions $F_{ij}^{(k)}(t) = 1 - e^{-\mu_i t}$, where μ_i are the state-dependent transition rates for the stochastic tree; and 4) three sets of rewards, corresponding to the amount of QALE expected in each state in X with respect to the decisions in Ω.

CHARACTERISTICS ANALYSIS

While sharing the general theoretical basis, the three dynamic decision modeling formalisms capture and reflect the underlying mathematical information in different ways; different assumptions are also adopted.

Time Indices and Horizons

The time index set T is discrete in dynamic influence diagrams and Markov cycle trees, and continuous in stochastic trees. The constant and exponential holding times captured in these models are very strong assumptions; they do not exploit the expressiveness of the flexible holding times in semi-Markov processes. For instance, facts like "after a PTCA, if the patient is well, it is three times more likely that he will develop restenosis in the next 6 months than later" cannot be easily captured in the existing frameworks; extensions such as "tolls" and additional "time slices" have to be incorporated.

The horizon of a dynamic decision problem can be finite or infinite. In finite horizon problems, the optimal policy is determined over a finite number of time periods N, *e.g.*, deciding on the best course of actions for 65 year-old patients in the 5 years immediately following a PTCA. In infinite horizon problems, the decision endpoints may be at an arbitrarily distant future, *e.g.*,- when a large cohort of patients die off, using a small cycle time unit. Since all the decisions and events have to be explicitly displayed in dynamic influence diagrams, they are feasible only for modeling finite horizon problems. Markov cycle trees can model both finite and infinite horizon problems. In stochastic trees, state transitions can occur at any time corresponding to some exponential distributions, hence they are valid only for infinite horizon problems.

Representational Explicitness

One major advantage of dynamic decision modeling is indeed the graphical modeling languages involved. Although different types of decision models explicitly display different types of information, any graphical tool would provide tremendous insights into the decision context for the decision maker.

Specifically, while the actual values of the chance events and decisions are concealed in the nodes, dynamic influence diagrams explicitly display the sequential nature of the decisions and probabilistic independencies among the variables. Independent health outcome variables can be dealt with separately. For instance, the dynamic influence diagram in Figure 2 explicitly shows that the health outcome variables *MI* and *Restenosis* are conditionally independent, and hence unnecessary to be combined into a single state variable as in the embedded Markov chain in Figure 1. The structure of the value function can also be explicitly displayed in these models.

As in conventional decision trees, the Markov cycle tree in Figure 3 explicitly displays the possible consequences of each state variable, which are implicitly captured in the transition arrows of the embedded Markov chain.The sequential nature of the decisions, however, cannot be captured in a straightforward manner. In fact, Figure 3 does not include sequential decisions, *i.e.*, it does not consider how backup and repeated treatments may affect the patient's prognosis. It is possible to incorporate such information, but the resulting Markov cycle tree would be too complicated to be displayed in the limited space here. The combinations of all subsequent actions and their possible consequences in one cycle would have to be modeled as separate states in the cycle tree; these subsequent decisions cannot be modeled as decision nodes in this framework.

The stochastic trees are simply the continuous analog of the Markov cycle trees, although only the immediately possible state transitions are depicted. The difficulties in incorporating sequential decisions carry over.

Strategic Constraints

In our example problem, all the alternate decisions are assumed to be applicable at every stage; the decisions are assumed to be independent. As mentioned earlier, only dynamic influence diagrams can explicitly and easily capture this sequential nature of decisions. In the tree-based formalisms, the same decision or a very limited, fixed sequence of decisions is assumed at every decision point, *i.e.*, they can only model problems that compare the effectiveness of different decisions, but not those that determine an optimal sequence of decisions.

Moreover, the independent decisions assumption may not be true in general. For example, the efficacy of PTCA might depend on the number and nature of the procedures that the patient had gone through before; a patient could only go through three CABG procedures; or a CABG procedure could only be followed by a PTCA or medical therapy.

Such strategic constraints can be explicitly represented in a dynamic influence diagram, but again cannot be easily incorporated into either a Markov cycle tree or a stochastic tree, unless one is willing to expand the tree structures fully before evaluation. Even in a dynamic influence diagram, all the constraints have to be anticipated at model building time, and hence fully displayed before evaluation.

Discount Factors

The values achievable in each state of the embedded Markov chain may be discounted over time. For instance, life-expectancy generally lowers with age. The discount factors in turn might be constant or time-dependent. Since all three modeling formalisms described so far involve explicit enumeration of the state transitions, incorporating discounting is straightforward.

Solution Techniques

Dynamic influence diagrams: The decision problem is solved by graph reduction. The chance nodes are removed by conditional expectations, the decision nodes by value maximization, and the non-terminal value nodes by merging into other value nodes, and ultimately into the terminal value node. The only well established algorithm for solving dynamic influence diagrams is due to Tatman and Shachter [12]; recent efforts have tried to reduce these models into Bayesian or probabilistic networks to take advantage of the much larger collection of evaluation algorithms [11], but all these algorithms are still NP-hard [3] with respect to the size of the models.

Markov cycle trees: The decision problem is solved by evaluating the Markov cycle trees at the end of the alternate treatments and compare their relative values. Evaluation of a Markov cycle tree is by simply rolling forward with an initial value assignment and probability distribution of the state variables at the root. The next cycle starts with a new distribution at the root of the cycle tree. This corresponds to a cohort analysis [1]. The expected values are accumulated for each state until the process converges, or ends when the "dead" state reaches probability 1; this is an infinite horizon problem. The evaluation may take an arbitrarily long time to terminate, or may not terminate at all.

Stochastic trees: Evaluation of the stochastic tree is by rolling back. The expected values of each state is calculated by value expectation in the subtree with the state in concerned as the root, along the branches from the leaves. If cycles appear in the state transition diagram of the embedded Markov chain, the evaluation must be done across several cycles until the numbers converge. Due to the decomposability of the exponential transition time functions, fewer calculations are required as compared to Markov cycle tree evaluation, and the evaluation is finitely terminating.

All the dynamic decision evaluation algorithms are based on the *value iteration* method of the dynamic programming or Bellman optimality equation of semi-Markov decision processes, the discrete time version of which is shown in EQ 1. This method is based on the notion of *backward induction*. All possible state evolutions over time have to be explicitly enumerated in the models, and hence considered by the evaluation algorithms.

Even then, the power of dynamic programming, *i.e.*, considering only the optimal decisions obtained so far from the future by backward induction, is actually not fully exploited in dynamic decision modeling. Although dynamic influence diagrams correspond to exact formulations of the dynamic programming equation [12], all the tree-based formalisms do not make use of the suboptimality structure. In these latter frameworks, decision nodes are not part of the model structures, hence all possible decision consequences have to be simulated during evaluation.

DISCUSSION

Regardless of the particular framework employed, dynamic decision modeling has proven to be quite formidable in many real applications, even with the help of computers. The computational complexity arises from the value iteration technique of the underlying dynamic programming formulations; all the relevant events over time must also be explicitly depicted in most cases. In order to circumvent such complexity, either strong strategic constraints that limit the number of decision choices at each decision stage are imposed, or approximate algorithms for evaluating partial models are employed [10] [5]. While these latter myopic strategies could possibly guide the evolution of the models more accurately, they do not guarantee global optimality of the solutions. In fact, no formal characterization of such strategies, in terms of the decision factors involved and the corresponding solutions, have been attempted.

Research in stochastic control and Markovian decision problems has resulted in a number of solution methods for semi-Markov decision processes [2] [9]. The applicability and efficiency of the different methods, however, depend on specific characteristics of the problems under consideration. For problems with finite decision and state spaces, certain discount factors, and stationary policies that involve decisions that are independent of time, methods such as *policy iteration*, *adaptive aggregation*, and *linear programming* [2] may be more efficient than value iteration. We believe, although we will not elaborate here, that the assumptions required can be satisfied by a substantial subclass of challenging clinical decision problems

Therefore, it seems that much complexity in dynamic decision modeling can be avoided by dealing directly with its underlying theoretical framework as a semi-Markov decision process. In particular, many problems that the models aim to solve can be solved by more efficient techniques. This conclusion provides some interesting insights into integrating graphical and numerical approaches to automatically solve dynamic decision problems in medicine.

CONCLUSION

Dynamic decision making is a knowledge- and labor-intensive task. Automating the decision making process has recently become an important area of research [13]. One major controversial issue in automated dynamic decision modeling research is the choice of the models employed, since different models adopt different assumptions, explicitly display different types of information, and hence would ease different types of sensitivity analysis to different extents. On the other hand, most direct solution techniques for semi-Markov decision processes do not take advantage of the graphical insights provided by the decision analytic approach.

Since all dynamic decision modeling formalisms are based on the same theoretical framework of semi-Markov decision processes, the same types of information are actually required to formulate a decision problem in both approaches. Only after the problem formulation process is completed, *e.g.*, all the decision alternatives, state variables, transition probabilities, *etc.*, have been identified, that the information apparent or easily accessible to the decision maker appears different. If all the information leading to the problem formulation can be traced somehow, the two approaches should not differ.

In conclusion, we believe that a feasible approach to automated dynamic decision making is to formulate the problems directly as semi-Markov decision processes. With a well-designed interface, the relevant information for problem formulation could be input and accessed in whatever convenient forms, graphical or otherwise. Only after the problem has been formally and completely formulated, a particular format is chosen, depending on the problem nature, to solve for the optimal policy. Sensitivity analysis can be conducted by accessing the recorded information, and perhaps also by employing a different solution technique. This hypothesis will be tested in a future project.

Acknowledgments

The author would like to thank Drs. Stephen G. Pauker and Michael P. Wellman for many helpful discussions. I am solely responsible for any error in the presentation of the ideas. This research is supported by the National Institute of Health grant no. 5 R01 LM04493 from the National Library of Medicine.

References

[1] J. Robert Beck and Stephen G. Pauker. The Markov process in medical prognosis. *Medical Decision Making*, 3:419–458, 1983.

[2] Dimitri P. Bertsekas. *Dynamic Programming: Deterministic and Stochastic Models*. Prentice-Hall, Englewood Cliffs, NJ, 1987.

[3] Gregory F. Cooper. The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence*, 42(2-3):393–405, 1990.

[4] Robert Gallager. Lecture notes on discrete stochastic processes. MIT Course Notes, 1992.

[5] Robert P. Goldman and John S. Breese. Integrating model construction and evaluation. In Didier. Dubois, Michael P. Wellman, Bruce D'Ambrosio, and Phillipe Smets, editors, *Uncertainty in Artificial Intelligence: Proceedings of the Eighth Conference*, pages 104–111, San Mateo, CA, 1992. Morgan Kaufmann.

[6] Gordon B. Hazen. Stochastic trees: A new technique for temporal medical decision modeling. *Medical Decision Making*, 12:163–178, 1992.

[7] Daniel P. Heyman and Matthew J. Sobel. *Stochastic Models in Operations Research: Stochastic Optimization*, volume 2. McGraw-Hill, 1984.

[8] J. P. Hollenberg. Markov cycle trees: A new representation for complex Markov processes. *Medical Decision Making*, 4(4), 1984. Abstract from the Sixth Annual Meeting of the Society for Medical Decision Making.

[9] Ronald A. Howard. *Dynamic Probabilistic Systems*, volume 1 & 2. John Wiley and Sons, New York, 1971.

[10] Gregory M. Provan. Modeling the evolution of acute abdominal pain using temporal influence diagrams. Unpublished manuscript, 1992.

[11] Ross D. Shachter and Mark A. Peot. Decision making using probabilistic inference methods. In D. Dubois, M. P. Wellman, B. D. D'Ambrosio, and P. Smets, editors, *Uncertainty in Artificial Intelligence: Proceedings of the Eighth Conference*, pages 276–283. Morgan Kaufmann, 1992.

[12] Joseph A. Tatman and Ross D. Shachter. Dynamic programming and influence diagrams. *IEEE Transactions on Systems, Man, and Cybernetics*, 20(2):365–379, 1990.

[13] Michael P. Wellman, John S. Breese, and Robert P. Goldman. From knowledge bases to decision models. *The Knowledge Engineering Review*, 7(1):35–53, 1992.

[14] John B. Wong, Frank A. Sonnenberg, Deeb M. Salem, and Stephen G. Pauker. Myocardial revascularization for chronic stable angina: Analysis of the role of percutaneous transluminal coronary angioplasty based on data available in 1989. *Annals of Internal Medicine*, 113(11):852–871, 1990.